

S/N: 10/762,839

Reply to Office Action of: 06/29/2006

Amendments to the Drawings:

Please replace the drawings with the new ones enclosed along with this document. There are a total of eight drawings, duly marked.



said method further enables all or selected nodes to write data across the said shared storage capacity of all or a select set of nodes or server groups in said network system.

19. A method for the effective utilization of data storage capacity in a communicatively coupled network system using a storage architecture wherein:

the storage capacity on all or selected nodes in said network system is segmented into non-sharable and shared storage capacity based either on the percentage of 'available/unused' or 'total installed' storage capacity or as a specific amount of unused or installed storage capacity on the nodes;
said method further enables all or selected nodes to write data to a central file server or server group after crossing a predetermined storage capacity threshold, and;
said data written onto said central file server is striped across said shared storage space on said nodes in said network system.

20. A method of claim 19, where the data written to the central file server or server group is abstracted and displayed as a local file with a pointer to the actual location on the central file server or server group.
21. A method of claim 19, where the data written to shared storage is batched together into a sequential log and then divided into efficient larger blocks to optimally use the network and I/O subsystem bandwidth.
22. A method of claim 19, where the data written to shared storage is scheduled at off-peak system usage times based on statistics provided by network bandwidth meters and I/O subsystem capacity.
23. A method of claim 19, where the data is written either at a file level or block level.

Remarks

Reconsideration and reexamination of the above-identified patent application, as amended, are respectfully requested. Claims 4-23 are pending in this application upon entry of this Amendment. In this Amendment, the Applicant has cancelled claims 1-3; and added new claims 4-23. Of the pending claims, claims 4, 11, 18, and 19 are the only independent claims.

Claim Rejections – 35 U.S.C. § 102

In the Office Action mailed June 29, 2006, the Examiner rejected independent claim 1 under 35 U.S.C. § 102(b) as being anticipated by U.S. Patent No. 5,784,697 issued to Funk ("Funk"); rejected independent claims 2

and 3 under 35 U.S.C. § 102(e) as being anticipated by U.S. PG Publication No. 2003/0028592 A1 by Ooho et al. ("Ooho") and U.S. Patent No. 6,415,373 B1 issued to Peters ("Peters").

The Applicant respectfully traverses the rejection to the claims under 35 U.S.C. § 102(b) and U.S.C. § 102(e). The Applicant has written newly added independent claims 4, ~~18~~, and ~~19~~ to more clearly define over the cited patents.

1. Background of the Claimed Invention

In any computing environment, it is desirable to effectively use the existing hardware and software resources before making additional infrastructure investments. Large corporations, universities etc. normally tend to have network infrastructures with multiple, heterogeneous servers and decentralized data. These servers are often chosen by or within specific departments for their specific advantages in tackling particular problems. For instance, Unix servers are known for handling complex design and engineering assignments while certain other servers offer superior random file access characteristics (for database applications). It naturally follows that there would be installed storage capacity on these servers and on their individual clients and independent disks/storage media under the control of the particular server, hereinafter referred to as a 'Server Group' (SG).

However, with traditional methods, additional storage capacity is installed when required for a particular server (for example, when data stored on a node is reaching or has reached critical maximum levels) while a large amount of such data storage capacity that already exists on other server groups (SG), on the network, may remain untapped. This represents unnecessary expenditure for the enterprise as a whole, which may require additional unnecessary work assignments such as requesting quotations from vendors, vendor/product evaluation, commissioning of new equipment etc. Further, the inability to seamlessly allocate storage capacity on servers and server groups in different departments of an organization - as a ratio of total installed capacity or unused storage capacity inhibits efficient planning and implementation of an organization wide storage segmentation policy. This prevents true utilization of organization wide data storage resources and results in unnecessary expenditure to buy and maintain additional yet unnecessary storage resources while the ones already installed on the network remain unused.

The following example further illustrates this point. Example: A large organization uses separate server clusters or storage area networks (SAN) for (a) managing all emails and attachments, (b) for archiving all architectural plans and catalogs and (c) for storing copies of all video advertisements it has put out. Let us assume that the nodes (that make up the server cluster or SAN) handling Email have an installed storage capacity of a hundred storage units, but an actual usage of eighty storage units, which means that twenty units of storage remains unused. The nodes handling the catalogs and architectural plans have an installed capacity of

hundred storage units and actual usage of only twenty storage units – leaving eighty storage units unused. Finally, the nodes handling video ads have an installed storage capacity of fifty storage units and current usage of twenty storage units. Further, all the organizational computing resources are coupled together by a network. During instances, when the Email nodes receive a sudden influx of emails and attachments and the existing storage capacity is rapidly being used up, it would be ideal to use at least a certain portion of the excess unused/installed storage capacity on the 'catalog' nodes and 'video ads' nodes to spillover excess data beyond the installed capacity on the dedicated Email nodes, while still leaving aside storage capacity on those nodes for their own dedicated usage. This would prevent any stoppage of service when the email handling nodes would otherwise have run out of storage space. At the same time, it would also give the organization the ability and much needed flexibility to handle sudden increases in data storage capacity and crucial time to subsequently plan and implement the exact amount of additional data resources, increase efficiency (by increased utilization) of existing resources and an increased ROI (return on investment) on their data storage resources while preventing avoidable additional capital investments.

As such, what is needed is an innovative, flexible storage architecture that enables segmentation of storage capacity as needed into shared and dedicated (unshared) storage based either on the available storage capacity or total capacity on the selected nodes, while taking into account the bandwidth of the network and the I/O subsystem. The claimed invention satisfies this need.

2. The Claimed Invention

a. Overview of the Claimed Invention

The claimed invention, as set forth in newly added independent claims 4, 11, 18 and ~~20~~, is generally directed to providing a novel, scalable network file system that uses multiple servers, individual clients and/or storage devices to effectively and optimally utilize existing storage resources, provides for high data availability, high system performance and greater protection against data loss by replicating data with parity and takes network bandwidth and I/O subsystem usage limitations into consideration. Some salient features of this invention are:

1. The invention enables dynamic configuration of storage capacity in the computing environment and segmentation of the same into dedicated and shared storage capacity from a single console. Further, this storage segmentation may be done as a percentage of total storage capacity on that node or only a portion of the available storage space.
2. The invention uses a centralized file server directly in the path of a distributed file system wherein, the data that is written by the nodes to global storage is actually written onto the file server and the back-up copy of

this data is then striped onto the global storage on the several selected nodes whose storage capacity has been segmented and contains space for data to be written by other nodes in the network. However, the back-up copy for the data that exists in the dedicated segment on the nodes (for use by the nodes themselves) is, in fact, written onto the centralized file server. So, the file server actually acts as a host of original data and also as a back-up device.

3. From the above two points, it is clear that the 'global storage' that has been segmented out from the selected nodes is actually a space where back-ups of data is written and is not for actual data written by other nodes to be striped across in real time (which would not take network and I/O subsystem bandwidth into consideration). This is a fundamental difference with other network files systems and is an improvement over the same.

4. The data that is written (backed-up) by the centralized file server onto the 'global storage' space is batched together into a sequential log and divided into optimum block sizes based on the network and I/O subsystem bandwidth.

5. The invention is modular and allows for additional storage capacity to be added to the existing nodes in real time or for new nodes to be added altogether to the network within the ambit of the implemented global and nodal storage policies.

b. Newly Added Independent Claims 4, 11, 18 and ~~20~~ 19

The claimed invention, as recited in newly added independent claim 4, is a method for creating an operating a communicatively coupled network system wherein the segmentation of storage capacity into non-sharable and share storage takes place. This segmentation may be done based either on the percentage of 'available/unused' or 'total installed' storage capacity or as a specific amount of unused or installed storage capacity on the nodes. Data may then be written by these nodes (upon exhaustion of the storage capacity) to a centralized file server or server group, which in turn, backs-up the data by striping the same onto the shared storage capacity on all or select nodes in the network. Network bandwidth and I/O subsystem usage are taken into consideration.

Newly added independent claim 11 recites a system claim for achieving the procedure mentioned above.

The storage segmentation and data striping system described in newly added independent claim 18 does not use any central file server or server group. After segmentation into shared and non-sharable segments, data may be written by nodes (upon exhaustion of the storage capacity on individual nodes or in that server group) to the shared storage capacity of a select group of nodes called a 'stripe set' to effectively utilize the data storage capacity available within the network system.

a centralized file server or server group, which in turn, backs-up the data by striping the same onto the shared storage capacity on all or select nodes in the network.

Newly added independent claim 19 recites a method for creating an operating a communicatively coupled network system using a storage architecture wherein the segmentation of storage capacity on nodes into non-sharable and share storage takes place. This segmentation may be done based either on the percentage of 'available/unused' or 'total installed' storage capacity or as a specific amount of unused or installed storage capacity on the nodes. Data may then be written by these nodes to a centralized file server or server group, which in turn, backs-up the data by striping the same onto the shared storage capacity on all or select nodes in the network.

3. The Claimed Invention Compared to Funk, Ooho and Peters.

Funk teaches a method of assigning processes in a multiprocessing system having non-uniform memory access. Processes are first linked to pools of memory spaces in the logical main memory. The logical main memory in turn extends all through the local main memories of the nodes.

Each multiprocessing node is assigned a specific process based on whether or not the process is associated with a pool that uses the local main memory on that particular node. This way, Funk avoids the necessity to access the memory on other nodes (which takes up more time, as mentioned by Funk in col. 1, lines 57-59), which generally happens in multi-processing systems.

Several fundamental differences exist between the invention and the prior art of Funk. In fact, the invention and that of Funk are so far apart that they cannot even be compared. Funk teaches a method for improving the system performance of a multi-processing system, which actually uses the memory and not the hard disk's storage as taught by the invention. In col.1, lines 10-11, Funk says: 'in particular to an improved memory distribution mechanism'. Further, in col.1, lines 52-55 Funk says: 'that allows all processors to access any of the main storage in the system. The nodes share the same addressable main storage, which is distributed among the local main memories of the nodes'. This clarifies beyond any doubt that the main storage is actually the memory on the nodes.

This is further reinforced in col. 2, lines 52-54 which identify the storage architecture in Funk's teaching as a memory cluster: 'Non-Uniform Memory Access (NUMA) storage architecture, also known as Shared Memory Cluster (SMC) storage architecture.

A definition of 'memory' provided by TechTarget – an IT media company is enclosed below:

Memory is the electronic holding place for instructions and data that your computer's microprocessor can reach quickly. When your computer is in normal operation, its memory usually contains the main parts of the operating system and some or all of the application programs and related data that are being used. Memory is often used as a shorter synonym for random access memory (RAM). This kind of memory is located on one or more microchips that are physically close to the microprocessor in your computer. Most desktop and notebook computers sold today include at least 16 megabytes of RAM, and are upgradeable to include more. The more RAM you have, the less frequently the computer has to access instructions and data from the more slowly accessed hard disk form of storage.

Memory is sometimes distinguished from storage, or the physical medium that holds the much larger amounts of data that won't fit into RAM and may not be immediately needed there. Storage devices include hard disks, floppy disks, CD-ROM, and tape backup systems. The terms auxiliary storage, auxiliary memory, and secondary memory have also been used for this kind of data repository.

Additional kinds of integrated and quickly accessible memory are read-only memory (ROM), programmable ROM (PROM), and erasable programmable ROM (EPROM). These are used to keep special programs and data, such as the basic input/output system, that need to be in your computer all the time.

The current invention deals with the storage capacity on the nodes' storage devices and not with the memory. The data that is written into the memory in Funk's teaching is ephemeral in nature, is discarded by the system and is replaceable after the process handling is completed and is also wiped off when the system is switched off, unlike data written onto the storage devices (which are a nonvolatile computer readable medium) as taught by the current invention. The purpose of Funk's teachings is to improve performance in a multi-processing system whereas the purpose of the invention is to achieve optimum utilization of existing data storage capacity.

Ooho teaches a data back-up restoration system in a client-server type environment that takes into consideration the ownership rights at the time of backup in order to avoid enabling restoration of data by people who are not entitled to do so while enabling backup to any terminal (or node) even if the original terminal is broken. This is explained in paragraph 0005. In paragraph 0007, Ooho teaches that the back-up database in figures 1 and 7 holds the rights (or entitlement) information indicating the usage rights at the time of backup. This is reiterated in paragraph 0041, lines 14-18. In other words, users who subscribe to content in the content distributions system are identified to have these rights by the issuance of a 'license ticket' (LT). Information about these licenses tickets is backed up in the 'backup database'. See paragraph 0060, lines 14-16 and lines 25-29. In figure 7 also, there are multiple distribution centers (see 720 and 721) and not just one.

The backup database as taught by Ooho is neither a central repository on a centralized server to hold a copy of all data on all nodes on the network as taught by the current invention nor does it even hold the entire data from even a single node (see paragraph 0051, all lines).


The teachings of Peters talk about a method to store multiple copies of data on multiple storage units in a random or pseudorandom fashion to enable its quick retrieval and transfer. This is directed more towards ease of reading and transferring high bandwidth streams of data when different applications request the same - see col. 2, lines 13-27. The current invention is fundamentally different from the teachings of Peter. Moreover, the claims in the invention have now been modified to further differentiate from Peters.

Accordingly, the Applicant believes that the claimed invention as recited in newly added independent claims 4, 11, 18, and 19 is patentable over Funk, Ooho and Peters. Claims 5-10, 12-17 and 20-23 depend from one of the independent claims and include all of the limitations therein. Thus, the Applicant respectfully requests reconsideration and withdraw of the rejection to the claims under 35 U.S.C. § 102(b) and under 35 U.S.C. § 102(e) in view of Funk, Ooho and Peters and also other prior art that was made of record and not relied upon.

CONCLUSION

In summary, newly added claims 4-23 meet the substantive requirements for patentability. The case is in appropriate condition for allowance. Accordingly, such action is respectfully requested. If a telephone or video conference would expedite allowance or resolve any further questions, such a conference is invited at the convenience of the Examiner.

Respectfully submitted,



SRIKRISHNA TALLURI

Date: November 27, 2006